Risk and Return in the Design of Environmental Policy

Robert S. Pindyck

Abstract: I examine risk/return trade-offs for environmental investments and their implications for policy choice. Consider a policy to reduce carbon emissions. To what extent should the policy objective be a reduction in the expected temperature increase versus a reduction in risk? Using a simple model of a stock externality that evolves stochastically, I examine the "willingness to pay" (WTP) for alternative policies that would reduce expected damages versus the variance of those damages. I compute "iso-WTP" curves (social indifference curves) for combinations of risk and expected return as policy objectives. Given cost estimates for reducing risk and increasing expected returns, one can compute the optimal risk-return mix for a policy, and the policy's social surplus. I illustrate these results by calibrating the model to data for global warming.

JEL Codes: D81, Q5, Q54

Keywords: Climate change, Economic growth, Environmental policy, Policy objectives, Risk, Risk-return trade-off, Uncertainty, Willingness to pay

THE ADOPTION of an environmental policy usually imposes costs on society, but it is expected to yield a social return in the form of a stream of benefits, for example, health benefits from less pollution. Those benefits, however, usually occur in the future and may be highly uncertain, so that (as with other private or public investments) the actual return is uncertain. This is especially true for environmental policies involving stock externalities, such as increases in greenhouse gas (GHG) concentrations, the acidification of lakes and oceans, and the accumulation of toxic waste.

Robert S. Pindyck is at Massachusetts Institute of Technology, Cambridge, MA 02142. My thanks to Stacie Cho, Lorna Omondi, and Andrew Yoon for their excellent research assistance, and to Ravi Bansal, John Cox, Richard Schmalensee, seminar participants at Brandeis, MIT, University of Quebec at Montreal, and Tel-Aviv University, and three referees for helpful comments and suggestions.

Received November 24, 2013; Accepted June 27, 2014; Published online September 2, 2014.

JAERE, volume 1, number 3. 2333-5955/2014/0103-00XX $10.00 http://dx.doi.org/10.1086/677949

All of these problems involve long time horizons and uncertain impacts. This paper examines the risk/return trade-off for environmental investments and the implications of that trade-off for policy design. To illustrate the basic idea, I utilize a very simple and stylized model of climate change, which involves a long time horizon and considerable uncertainty. However, the framework developed here could well be applied to other environmental (and nonenvironmental) policy problems, such as the management of toxic waste and the development and protection of water resources.[1] My use of climate change is done purely as an illustrative example.

Consider a policy to reduce GHG accumulation, for example, a carbon tax. The policy would have some cost to society, but by reducing GHG emissions it would reduce the extent of warming by some uncertain amount and thus yield uncertain future benefits. The uncertainty arises because we have limited knowledge of the relationship between GHG concentrations and temperature, and between higher temperatures and GDP growth. A question then arises. To what extent is the social value of the policy (measured by society's willingness to pay for it) driven by its expected benefits versus uncertainty over those benefits?

The answer to this question can guide policy design. Alternative policies might have different impacts on the expected change in temperature over the next century versus the variance of that change. For example, rather than (or in addition to) GHG emissions abatement, money might be spent to improve our knowledge of how GHG emissions affect temperature and how changes in temperature affect economic output. A purely research-focused policy might do little or nothing to reduce expected future damages, but it might reduce the variance of those damages. Likewise, investments in adaptation to climate change (e.g., the development of hardier hybrid crops) could reduce the right tail of the distribution of future damages, reducing both the variance and expectation of damages.[2]

There is value in reducing the expected damages from GHG emissions and in reducing the variance of those damages; at issue is the trade-off between the two. To what extent should policy be aimed at reducing the expectation versus the variance of damages?[3]

The framework I use to address this question is "willingness to pay" (WTP). Consider a policy that would reduce expected future temperature increases and/or

1. For a general discussion of the nature of the uncertainties involved in environmental policies, whether or not stock externalities are involved, see Pindyck (2007). This paper is a shorter version of Pindyck (2012a).

2. This would also affect the skewness of the distribution of future damages. As discussed below, I focus on variance as the measure of uncertainty and ignore higher moments.

3. A separate question is how learning (and thus reductions in variance) that occurs naturally as the climate system evolves should affect a policy to reduce expected damages. See, e.g., Kelly and Kolstad (1999).

the variance of those increases. The WTP for the policy is the maximum percentage of current and future consumption that society would give up to achieve those reductions. This does not mean it would be possible to achieve these reductions given the revenues from this WTP; doing so may be more costly, or less costly. WTP relates only to the demand side of policy: it is society's reservation price for the achievement of particular policy objectives.

In Pindyck (2011b, 2012b), I calculated the WTP for GHG abatement policies that would limit future increases in temperature to some specified amount (e.g., 3°C), based on a probability distribution for the rate of warming under "business as usual" (BAU). Of course WTP depends on the probability distribution, and in particular on both the expected temperature increase under BAU and the variance of the increase. I calculated the trade-off between the expected value of a temperature increase versus its standard deviation under BAU as determinants of WTP. That trade-off is the marginal rate of substitution along an "iso-WTP curve," that is, the locus of combinations of the expected value versus the standard deviation of the change in temperature under BAU such that the WTP is constant. The results showed that for a society that is risk averse, uncertainty over temperature change can indeed be a strong driver of WTP.

In that earlier work, I took the policy objective—limiting the temperature increase to some amount—as given and simply examined the implications of the starting conditions, that is, the extent to which the WTP for the policy depends on the variance of the temperature increase. Here I turn to the design of policy by considering alternative objectives and, in particular, the trade-off between reducing the expected rate of increase in temperature versus reducing its variance. Thus I ask, What combinations of drift reduction and variance reduction yield the same WTP? Those combinations trace out a social indifference curve—combinations of changes in drift and variance that are welfare equivalent. This social indifference curve provides a framework for policy design in that any policy objectives on the curve have the same impact on welfare. To my knowledge, this paper is the first to address environmental policy design in the context of a trade-off between risk and return.

Given this social indifference curve, what combination of risk and return should society aim for? If one could estimate the costs of drift reduction and variance reduction (no easy task), we could take this a step further and compare the cost of a particular combination of drift and variance reduction to the WTP for that combination, that is, whether the policy yields a positive social surplus. Finally, assuming the costs are weakly convex (as we would expect), we can determine the optimal risk-return trade-off for policy design.

Why focus only on the variance of changes in an environmental variable (such as temperature) and not also on higher-order moments (such as skewness and kurtosis)? First, I want a model that is reasonably simple and easy to understand, both in terms of the process for the environmental variable and the utility function used to measure welfare. Second, I want to draw an analogy to the mean-variance risk-return trade-off facing investors in financial markets, an analogy that I believe provides insight and also creates a bridge between finance and environmental economics. Finally, expanding the social utility function so that it "values" higher-order moments requires ad hoc assumptions that are hard to justify.

In the next section, I lay out a simple model of a stock externality in which an environmental variable evolves stochastically and affects the growth rate of consumption. Using a simple constant relative risk aversion (CRRA) social utility function, I show how the model can be used to calculate the WTP for policies to reduce expected damages and/or the variance of damages, and I show how one can find the optimal target mix. In section 2, as an illustrative exercise, I do a rough calibration of the model to information on global warming and its impact, and then show how the model can be used to address risk-return trade-offs in the design of policy. Section 3 concludes.

1. RISK, RETURN, AND WILLINGNESS TO PAY

To illustrate the idea of a risk-return policy trade-off, I assume that there is an environmental stock variable X_t which is "bad," in that it reduces the growth rate g_t of GDP and consumption. Under BAU, X_t follows an arithmetic Brownian motion (ABM) with drift and volatility α_X and σ_X, respectively. A (costly) environmental policy could be designed and implemented to reduce α_X and/or σ_X, and thereby reduce the expectation and/or variance of future reductions in consumption that would otherwise result from the growth of X_t. Welfare is measured using a CRRA utility function, discounted at rate δ. Thus welfare at time t is

$$U(C_t) = C_t^{1-\eta}/(1 - \eta), \tag{1}$$

where η is the index of relative risk aversion (and $1/\eta$ is the elasticity of intertemporal substitution).[4]

4. Why use a simple CRRA utility function and not something more general? For example, one might introduce a continuous-time version of Epstein-Zin recursive preferences (e.g., as used in Pindyck and Wang [2013]) which would unlink the index of risk aversion from the elasticity of intertemporal substitution, or some other aspect of consumer preferences such as prudence (which would make the skewness of the distribution of future consumption particularly important). One reason for using a CRRA utility is that I am putting a premium on making the model simple and clear; recursive preferences, for example, would add considerable complexity with no equivalent gain in economic insight. Second, I am looking at the preferences of a policy maker and not the preferences of the individual consumers in the economy. As is well known (see, e.g., Samuelson 1956), without very restrictive and unrealistic assumptions, one cannot derive a social utility function by aggre-

What about the environmental stock variable, X_t? To put this in the context of a concrete example, I will assume that X_t is the anthropomorphic increase in temperature from its current level.[5] The process for X_t is

$$dX_t = \alpha_X dt + \sigma_X dz, \tag{2}$$

where dz is the increment of a Weiner process, and $X_0 = 0$. Note that equation (2) is essentially a vehicle for describing the expectation and variance of the temperature change from $t = 0$ to any future point in time T. In particular, $\varepsilon_0(X_T) = \alpha_X T$ and $\mathrm{Var}(X_T) = \sigma_X^2 T$.

The assumption that X_t follows an arithmetic Brownian motion is made for analytical convenience. It implies that both $\varepsilon_0(X_T)$ and $\mathrm{Var}(X_T)$ are unbounded as T grows, which one could argue is unrealistic. Of course I could have allowed X_t to follow a more general and complex process, for example, a mean-reverting process or a process with reversion to a stochastic trend (which would bound $\varepsilon_0(X_T)$ and $\mathrm{Var}(X_T)$ for any T). However this would greatly complicate matters and preclude an analytical solution of the model, without providing any additional economic insight. Also, I can obtain solutions for an arbitrarily large (but finite) value for T, for example, 300 or 400 years.

The growth rate of consumption is given by a simple linear relationship:[6]

$$g_t = g_0 - \gamma X_t, \tag{3}$$

where g_0 is the growth rate absent the stock variable, that is, when $X_t = 0$. Thus g_t also follows an ABM:

---

gating individual utility functions. Thus even if consumers have, say, recursive preferences, there is no reason to expect such preferences to apply to an individual or group designing a social policy.

5. In the context of global warming, the actual stock variable is the atmospheric GHG concentration, which in turn drives temperature change, but with a lag. I ignore the lag and thus can treat temperature change itself as the stock variable. In the context of toxic waste, the stock variable could be the quantity or concentration of waste material over some geographic area.

6. Most integrated assessment models relate the temperature increase T to GDP through a "loss function" $L(T)$, with $L(0) = 1$ and $L' < 0$, so GDP at a horizon H is $L(T_H)\ \mathrm{GDP}_H$, where GDP_H is but-for GDP with no warming. (Nordhaus [2008], e.g., uses an inverse-quadratic function.) The loss function $L(T)$ implies that if temperatures rise but later fall, GDP could return to its but-for path with no permanent loss. Theoretical arguments and empirical evidence support the view that higher temperatures (and environmental damage in general) should affect the growth rate of consumption rather than its level, as in the model I use here. For an analysis of the policy implications of a direct versus growth rate impact, see Pindyck (2011b).

$$dg = -\gamma\alpha_X dt - \gamma\sigma_X dz \equiv -\alpha dt - \sigma dz. \tag{4}$$

Once again, I am keeping the model as simple as possible. For example, some environmental shocks (including accumulated nuclear waste, sulfur emissions, and global warming) are likely to affect both the growth rate and the level of consumption, particularly if consumption is measured in quality-adjusted terms. Likewise, we might expect the growth rate g_t to be subject to both temporary and permanent shocks, as in the "long-run risk" model of Bansal and Yaron (2004), which has become widely used in the asset-pricing literature. The model could be generalized along these lines, but at the cost of complicating the analysis.

Equation (3) implies that at any time s,

$$g(s) = g_0 - \alpha s - \sigma\int_0^s dz = g_0 - \alpha s - \sigma z(s),$$

so consumption at a future time t can be written as

$$C_t = C_0 e^{\int_0^t g(s)ds} = C_0 \exp\left[g_o t - (1/2)\alpha t^2 - \sigma\int_0^t z(s)ds\right], \tag{5}$$

Given the CRRA utility function, welfare (under business as usual) at time 0 is then

$$W_0 = \frac{1}{1-\eta}E_0\int_0^\infty C_t^{1-\eta}e^{-\delta(t-s)}dt, \tag{6}$$

where δ is the rate of time preference, that is, the rate at which utility is discounted.

1.1. Willingness to Pay

Suppose that with some expenditure over time, society could reduce α_X and/or σ_X, that is, reduce the expectation and variance of future temperature change. Let α_X and σ_X be the drift and volatility of X_t under BAU, and α'_X and σ'_X be the corresponding drift and volatility under a policy that has a permanent cost to society of w percent of consumption. WTP is the maximum value of w that society would accept to achieve $(\alpha_X, \sigma_X) \to (\alpha'_X, \sigma'_X)$.

While designing a policy that reduces α_X might seem straightforward, how as a practical matter could a policy be designed to reduce σ_X, that is, to reduce the variance of the stock variable T years from now? Investment in research is the most obvious example. As I discuss later in the context of climate change, research might provide better estimates of the rate of temperature increase and better estimates of the impact of higher temperatures. Likewise, research in marine biology and population ecology could reduce uncertainty over the future damages to fisheries from increasing levels of acidity in lakes and oceans. Investments in adaptation could also reduce uncertainty over future damages from environmental degradation. A (perhaps extreme) example is the development of geoengineering technologies that could limit or even reverse temperature increases.[7]

WTP is society's reservation price for achieving the policy objective $(\alpha_X, \sigma_X) \to (\alpha'_X, \sigma'_X)$, but it is a reservation price for a particular type of policy and a particular form of pricing: at time $t = 0$ the policy is adopted (with no option to wait for more information), and the drift and volatility of X_t are immediately changed to their new values. The payment flow starts at $t = 0$ and continues forever, that is, takes the form of a permanent reduction of consumption of w percent (yielding a permanent flow of revenue to pay for the policy objective). This is probably unrealistic. First, we would expect that changing α_X and/or σ_X would take time. Second, there is little reason for the flow cost of the policy to be a fixed percentage of GDP or consumption. Thus WTP might instead be defined as the willingness to give up a particular time-varying percentage of consumption in return for a gradual shift in α_X and/or σ_X. However, the simplifications I impose add clarity to the basic results.

Welfare under the policy is

$$W_1(\alpha'_X, \sigma'_X) = \frac{(1-w)^{1-\eta}}{1-\eta} \mathcal{E}'_0 \int_0^\infty C_t^{1-\eta} e^{-\delta t} dt, \tag{7}$$

where $\mathcal{E}'_0$ denotes the expectation at $t = 0$ when the drift and volatility of X_t are α'_X and σ'_X. Under BAU (i.e., no policy), welfare is

$$W_2 = \frac{1}{1-\eta} \mathcal{E}_0 \int_0^\infty C_t^{1-\eta} e^{-\delta t} dt, \tag{8}$$

where $\mathcal{E}_0$ is the expectation under the original drift and volatility, α_X and σ_X. Then WTP is the value w^* that equates W_1 and W_2.[8]

From equations (3) and (4), the drift and volatility of the stock variable X_t correspond to a drift and volatility of the real growth rate g_t, that is, $\alpha = \gamma\alpha_X$, $\sigma = \gamma\sigma_X$, $\alpha' = \gamma\alpha'_X$, and $\sigma' = \gamma\sigma'_X$. In what follows I will refer to changes in α and σ rather than α_X and σ_X.

---

7. Geoengineering might involve seeding the atmosphere with sulfur particles. For a discussion of this and other geoengineering approaches to climate change, see Barrett (2008, 2009).

8. I treat GDP and consumption as interchangeable and assume that all losses from higher temperatures, including health effects and ecosystem damage, can be monetized and included in GDP.

1.2. Iso-WTP Curves

Suppose that under BAU, the drift and volatility of g_t are α_0 and σ_0, respectively. If we specify a target drift and volatility, α' and σ', we can calculate the WTP for a policy to achieve this target. Suppose that WTP is w_1^*. We are interested, however, in the trade-off between drift reduction and variance reduction as targets of policy. That is, we want to know what combinations of drift reduction and variance reduction will yield the same $\text{WTP} = w_1^*$, assuming we start at α_0 and σ_0. We can compute the locus of such combinations and thereby obtain an "iso-WTP curve," that is, combinations of α and σ which, starting at α_0 and σ_0, all have the same WTP and thus are welfare equivalent.

A hypothetical iso-WTP curve is illustrated in figure 1. Point *A* is the starting drift and volatility, and points *B* and *C* are two targets, both of which have the same WTP of $w_1^* = .03$. Moving from *A* to any point on the curve will have the same WTP of .03. Note that α increases as we move down the vertical axis, so the curve represents a social indifference curve between expected return (a lower value of α) and risk (a higher value of σ). It is thus analogous to the return-risk indifference curve of an investor choosing a portfolio of equities and a risk-free asset. The figure also shows another iso-WTP curve, labeled WTP = .05. It traces out combinations of α and σ that are lower and thus harder to attain. Moving from α_0 and σ_0 to any point on this curve would have a WTP of .05.[9]

If we knew the costs of reducing α and σ, we could plot an iso-cost line. If those costs were linear in the change in α and the change in σ, the iso-cost line would be a straight line, as drawn in figure 1. Its tangency with the iso-WTP curve $w_1^* = .03$ at point *B* is the cost-minimizing target combination of α and σ. There are other target combinations of α and σ that have a WTP of .03, but they would be more costly to achieve. Thus, point *B* represents the optimal risk-return policy target consistent with a WTP of .03.

Environmental policy design is usually (always?) based on expected benefits and costs. But figure 1 shows that a "mean-only" policy target is inefficient. An optimal policy target is likely to include a reduction in risk as well as a reduction in expected damages.

In figure 1, moving from point *A* to point *B* has the same WTP = .03 as moving from point *A* to point *C*. We can also calculate the combination of starting values for α_0 and σ_0 that yield the same WTP = .03 when moving to point *B*. In the case of

9. Figure 1 suggests that environmental uncertainty carries a significant welfare cost. Lucas (1987) has shown that the welfare cost of the uncertainty affecting the growth rate of consumption is negligible. However, that finding results from the assumption that stochastic shocks to the growth rate of consumption are mostly temporary, whereas here an increase in temperature has a permanent impact on consumption. As Martin (2008) has shown, permanent shocks imply a very high welfare cost.

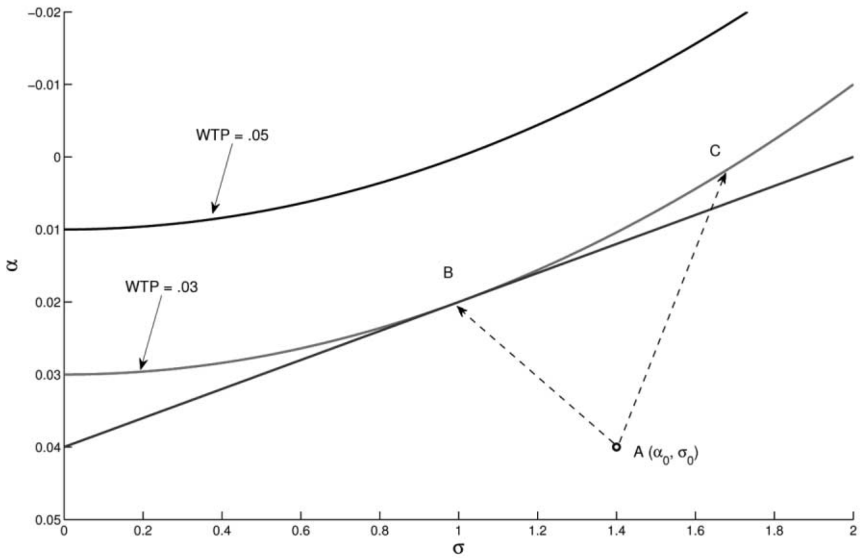

Figure 1. Iso-WTP curves and iso-cost line

global warming, for example, doing so is useful because it allows us to determine combinations of temperature drift and volatility under BAU that have the same welfare implications. In other words, how important is the expected change in temperature (under BAU) relative to its variance as determinants of the demand for policy? This is essentially what I did in Pindyck (2012b), that is, I focused on combinations of starting values for α and σ rather than target values.

This is illustrated in figure 2, which shows two iso-WTP curves. Curve 1 gives combinations of starting values α_0 and σ_0 such that the movement to point B has a WTP = .03. Thus, $w^*(A \rightarrow B) = w^*(D \rightarrow B) = w^*(A \rightarrow C) = .03$. But $w^*(D \rightarrow A) = w^*(B \rightarrow C) = 0$, so $w^*(D \rightarrow C) = .03$. Thus, moving from any point on curve 1 to any point on curve 2 will have the same WTP.

The iso-WTP and iso-cost curves in figures 1 and 2 are just illustrative. Later I will provide a very rough calibration of this simple model to estimates of temperature change and confidence intervals compiled by the IPCC (2007a, 2007b, 2007c) and others, and then calculate iso-WTP and iso-cost curves along the lines of those shown in figures 1 and 2.

1.3. Expected Utility

As can be seen from equations (7) and (8), in order to compute WTP we will need the expectation of the utility of consumption at future points in time, for specified values of α and σ. The appendix derives a formula for the expectation of $C_t^{1-\eta}$ for $t > 0$. Using that formula, the expectations in equations (7) and (8) can be written as

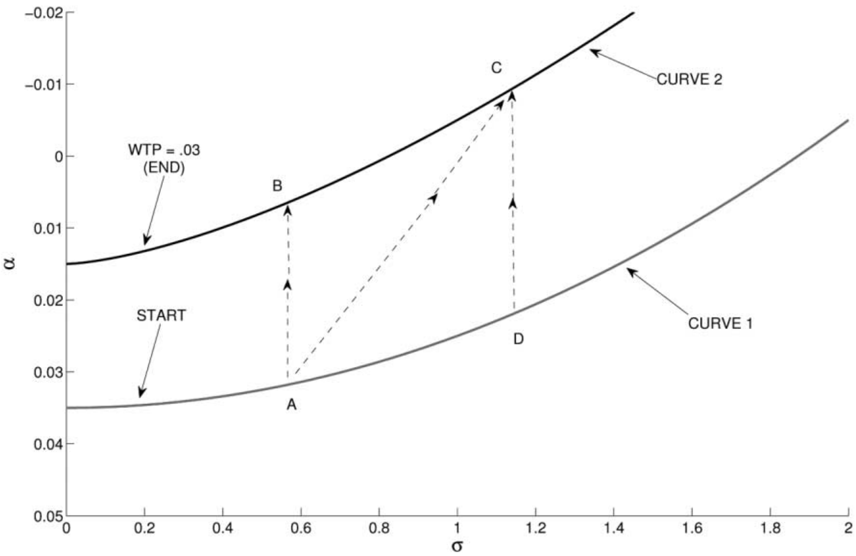

Figure 2. Iso-WTP starting curve and target curve

$$\mathcal{E}_0(C_t^{1-\eta})e^{-\delta t} = e^{-\rho_0 t + a(\alpha,\sigma,t)}, \tag{9}$$

where

$$\rho_0 \equiv \delta + (\eta - 1)g_0,$$

and

$$a(\alpha, \sigma, t) \equiv -(1/2)\alpha(1-\eta)t^2 + (1/6)\sigma^2(1-\eta)^2 t^3.$$

The σ^2 term in $a(\alpha, \sigma, t)$ is unbounded as t increases, so the integration in equations (7) and (8) must be done over a finite horizon. As mentioned earlier, if X_t (and thus g_t) followed a mean-reverting process, means and variances would be bounded, but that would greatly complicate matters and require numerical solutions. When calculating WTP, I limit the horizon to $T = 300$ years. It is hard to even imagine what the world might look like more than 300 years from now, and in any case, the results I present are intended to illustrate an approach to policy design rather than to evaluate any specific policy.

1.4. WTP and Policy Targets

We can now examine the WTP for alternative policies. Welfare over some future time horizon T with no policy intervention is given by

$$W_2 = \frac{1}{1-\eta}\int_0^T e^{-\rho_0 t + a(\alpha_0,\sigma_0,t)}dt. \tag{10}$$

Consider a policy to move from (α_0, σ_0) to (α_1, σ_1), and denote its WTP by w_1. With the policy, welfare is

$$W_1 = \frac{1}{1-\eta}\int_0^T (1-w_1)^{1-\eta} e^{-\rho_0 t + a(\alpha_1,\sigma_1,t)}dt. \tag{11}$$

Equating W_1 and W_2 yields

$$w_1 = 1 - \left[\frac{G(\alpha_1, \sigma_1)}{G(\alpha_0, \sigma_0)}\right]^{1/(\eta-1)}, \tag{12}$$

where

$$G(\alpha_0, \sigma_0) = \int_0^T e^{-\rho_0 t + a(\alpha_0,\sigma_0,t)}dt, \tag{13}$$

and likewise for $G(\alpha_1, \sigma_1)$. Thus, given starting values α_0 and σ_0 (based, say, on a calibration such as that discussed in the next section), we can calculate the WTP to change α and σ.

To obtain iso-WTP curves, that is, combinations of α' and σ' for which the WTP $= w_1$, we find the values of α' and σ' that satisfy

$$G(\alpha', \sigma') = G(\alpha_1, \sigma_1). \tag{14}$$

We can also obtain combinations of α' and σ' for which the WTP is equal to some arbitrary number, w. From equation (12), those combinations satisfy

$$G(\alpha', \sigma') = (1-w)^{\eta-1}G(\alpha_0, \sigma_0). \tag{15}$$

1.5. Costs of Changing α and σ

We have seen how to calculate the WTP for moving from (α_0, σ_0) to a new (α', σ'), as well as combinations of α and σ targets that have the same WTP. It is also useful to know the cost of moving from (α_0, σ_0) to (α', σ'). If that cost exceeds the WTP, the policy is economically infeasible; if the cost is less than the WTP, the policy yields a positive surplus.

Given the costs of changing α and σ, one can plot an iso-cost line, which would be a straight line (as in fig. 1) if the costs were linear in the changes in α and σ. Its tangency with an iso-WTP curve would be the cost-minimizing target combination of α and σ. The minimum cost, however, might be greater or less than the particular WTP. Given a starting point (α_0, σ_0), the WTP to move to a target (α', σ') depends on the distance to the target and also on the parameters η, δ, and g_0. For any given shift in α and σ, for example, a reduction in η will increase the WTP, which implies that for any fixed WTP, the required change in α and σ will be smaller, so that the corresponding cost will be smaller.

2. ILLUSTRATIVE EXAMPLE: CLIMATE CHANGE

To show how this framework could be applied to policy design, I determine values of the parameters that are roughly consistent with recent studies of global warming and its impact. I then calculate iso-WTP curves and discuss their implications. The model, of course, is extremely simple, so this exercise should be viewed as no more than an illustrative example. In fact, I have argued elsewhere—see Pindyck (2013a)—that integrated assessment models (IAMs) are of little or no value as tools for policy analysis, and the model used here is too simple to even be considered a "stripped-down" version of an IAM. The model is intended to illustrate the basic concepts presented in this paper and not to evaluate any specific policy. However, this exercise does show that a "mean-only" policy target is likely to be inefficient.

2.1. Parameter Values

To obtain numerical results we need values for α and σ in equation (3), the parameters η, δ, and the initial growth rate g_0, and estimates of the costs of reducing α and σ.

Calibration of α and σ under BAU

The 2007 IPCC report states that growing GHG emissions (i.e., under BAU) would likely lead to a doubling of the atmospheric CO_2-eq concentration relative to the pre-industrial level by mid-century, which would "most likely" cause an increase in global mean temperature between 2.0°C and 4.5°C by 2100, with an expected value of 2.5°C to 3.0°C. The IPCC indicates that this range, derived from the results of 22 scientific studies it surveyed, represents a roughly 66%–90% confidence interval, that is, there is a 5% – 17% probability of a temperature increase above 4.5°C. The IPCC's summary of the 22 studies also suggests that there is a 5% probability that a doubling of the CO_2-eq concentration would lead to a temperature increase of 7°C or more. For a 100-year horizon, I will use $\mathcal{E}(X) = 3$°C and a 5% probability of a temperature increase ≥7°C. The 5% point is 1.65 standard deviations above the mean (because eq. [2] implies that X_t is normally distributed), so one standard deviation is $4/1.65 = 2.42$. Thus the drift and volatility of the process for temperature are $\alpha_X = 3/100 = .03$, and $\sigma_X = 2.42/\sqrt{100} = 0.242$.

We need an estimate of γ, which relates g_t to X_t. One way to estimate γ is to combine expected future consumption with estimates of the expected loss of GDP at a specific temperature change. The appendix shows that expected future consumption is given by

$$\mathcal{E}_0(C_t) = e^{g_0 t-(1/2)\alpha t^2+(1/6)\sigma^2 t^3}. \tag{16}$$

In earlier work (2011b, 2012b), I used the IPCC (2007a, 2007b, 2007c) estimates that the loss of GDP resulting from a temperature change of 4°C is "most likely" in the range of 1% – 5%. Taking the upper end of the range, that is, 5%, implies that with $\alpha'_X = .04$, that is, $\mathcal{E}_0(X_{100}|\alpha'_X = .04) = 4°\text{C}$, and using equation (16), for $t = 100$,

$$\mathcal{E}_0(C_t) = \exp\left[g_0 t - (1/2)\alpha'_X \gamma t^2 + (1/6)\sigma_X^2 \gamma^2 t^3\right] = .95 e^{g_0 t},$$

so that

$$-\frac{1}{2}\alpha'_X \gamma t^2 + \frac{1}{6}\sigma_X^2 \gamma^2 t^3 = ln(.9). \tag{17}$$

Substituting $t = 100$, $\alpha'_X = .04$, and $\sigma_X = .242$ yields the following quadratic equation for γ:

$$\gamma^2 - .0205\gamma + .00000526 = 0.$$

If $\sigma_X = 0$, $\gamma = .000152$, so we want the smaller root of this equation, which yields $\gamma = .00026$.

This is an indirect estimate of γ and relies on models surveyed by the IPCC that largely posit, rather than estimate, a relationship between temperature and GDP. An alternative is to obtain γ from direct econometric estimates. Recent estimates include those of Dell, Jones, and Olken (2009, 2012) and Bansal and Ochoa (2011, 2012). The latter estimates are based on a more extensive data set (147 countries over 1950–2007) and allow for dynamic adjustment of GDP to changes in temperature. Bansal and Ochoa find that a 0.2°C increase in temperature leads to about a 0.2 percentage point reduction in the growth of real GDP. However, this impact number is transitory; GDP growth largely recovers after 10 years. In my model, the impact is permanent, so the Bansal-Ochoa number would imply a value of γ in equation (3) of around .001. This is four times as large as the value obtained from the IPCC, so I will use a value in the middle, namely, $\gamma = .0005$.

Thus we have $\alpha_X = .03$, $\sigma_X = .242$, and $\gamma = .0005$, so that $\alpha = \alpha_X\gamma = .000015$ and $\sigma = \sigma_X\gamma = .00012$. These are the values for α and σ that correspond to the starting values (α_0, σ_0) when calculating WTP and iso-WTP curves.

Other Parameter Values

WTP also depends on the index of relative risk aversion η, the rate of time preference δ, and the base level growth rate g_0. The historical per capita real growth rate g_0 is about .02. The finance and macroeconomics literatures put δ between .02 and .05, but it has been argued that for intergenerational comparisons δ should be zero. Without taking sides in that argument, I set $\delta = 0$. Likewise, values of η above 4 are consistent with the behavior of investors and consumers, but we might apply lower values to welfare comparisons involving future generations.[10] I will generally set $\eta = 2$, but I will also show how the results depend on η.

Costs of Reducing α and σ

Estimating these costs is speculative at best, given the paucity of available data, but I will proceed on the grounds that some estimate is at least illustrative. I use a linear cost function, that is, the cost of moving from (α_0, σ_0) to (α', σ') is:

$$\text{Cost} = c_1(\alpha_0 - \alpha') + c_2(\sigma_0 - \sigma'). \tag{18}$$

Note that the total cost and its two components are expressed in terms of percentages of consumption that must be sacrificed annually. Thus $c_1(\alpha_0 - \alpha')$ is the percentage of consumption that must be given up each year in order to reduce α from α_0 to $\alpha' < \alpha_0$.

I am aware of no direct estimates of c_1 or c_2, so coming up with numbers requires extrapolation of the few cost studies that exist, along with some guesswork. I start with the Kyoto Protocol, designed to limit any increase in global mean temperature by the end of the century to roughly 3°C. The US Energy Information Administration (1998) estimated that compliance with the protocol would cost from 1%–3% of GDP annually. Estimates from country cost studies in IPCC (2007a, 2007b, 2007c) also put the cost of limiting the increase in temperature to 3°C at around 1%–3% of GDP. I take the mid-range of 2% of GDP and translate "limiting the increase in temperature to 3°C" to mean reducing the expected temperature increase from 3°C to 1.5°C with a standard error of about 1.5°C. This means cutting the drift α in half, that is, from .000015 to .0000075, so that $c_1 = (.02)/(7.5 \times 10^{-6}) = 2{,}667$.

Estimating c_2, the cost of reducing σ, is more difficult. According to Palmer (2011), US public funding for climate change research in 2009 was about \$5 billion. These expenditures could arguably reduce uncertainty over future temperature change; they exclude public R&D spending on renewable energy, new technologies, and electricity and fossil fuel production as well as (the very limited) R&D spending directed at adap-

10. As Dasgupta (2008) pointed out, in the (deterministic) Ramsey growth model, the optimal savings rate is $s^* = (R - \delta)/\eta R$, where R is the consumption discount rate and the real return on investment. If $R = .04$ and $\delta = .02$, $s^* = 1/2\eta$, suggesting that η should be in the range of 2–4. Stern (2007) used a value of 1 for η, which is below the consensus range. For a general discussion of discount rates, see Gollier (2013).

tation to climate change.[11] I will assume that a doubling of the annual $5 billion in expenditures to $10 billion could result in a halving of σ_0. In other words, reducing σ from .00012 to .00006 would require sacrificing about 0.067% of GDP annually (based on a $15 trillion GDP), which implies that $c_2 = .00067/.00006 = 11.17$. This number should be viewed as just a "guesstimate," but it is illustrative. I also examine the implications of alternative values for c_2.

2.2. Some Calculations

Figure 3 shows three iso-WTP curves, generated using the starting values discussed above, that is, $\alpha_0 = .000015$ and $\sigma_0 = .00012$. Other parameters are $g_0 = .02$, $\delta = 0$, $\eta = 2$, and the time horizon is $T = 300$ years. The starting values (α_0, σ_0) are shown as the small circle near the bottom of the diagram. The solid line in the middle is the iso-WTP curve for WTP = .0365, which is the WTP to reduce α to zero but leave σ unchanged. Thus moving from (α_0, σ_0) to any point on that curve has a WTP of .0365. The other two curves in figure 3 show target combinations of α and σ that, starting from α_0 and σ_0, have WTPs of .02 and .06. (Note that all of the values of α on the top curve are negative.)

Figure 4 shows an iso-WTP curve (labeled B) for the same starting and ending values of the drift and volatility, $(\alpha_0, \sigma_0) = (.000015, .00012)$ and $(\alpha', \sigma') = (0, .00012)$, and the same values for the other parameters, so that the WTP is again .0365. The curve labeled A shows alternative combinations of *starting values* α_0 and σ_0 that give the same WTP = .0365 for any target on curve B. Thus, shifts from any points on curve A to any points on curve B have a WTP of .0365 and are welfare equivalent.

The iso-WTP curves in figure 3 show the most society would pay (in terms of percentage of consumption sacrificed) to move from a starting point (α_0, σ_0) to an ending point on the curve. But they do not tell us anything about the cost of that movement. If the cost exceeds the WTP, the policy is economically infeasible; if the cost is less than the WTP, the policy has a positive social surplus. In addition, since all points on an iso-WTP curve are welfare equivalent, we might want to know which target point has the least cost. To address this I use the iso-cost lines implied by equation (18), with the calibrated values for c_1 and c_2, that is,

$$\text{Cost} = 2,667(\alpha_0 - \alpha') + 11.17(\sigma_0 - \sigma'). \tag{19}$$

Figure 5 shows an iso-WTP curve that plots combinations of target values of α and σ that have a WTP of .02, given the starting point $(\alpha_0, \sigma_0) = (.000015, .00012)$,

11. For the six focus areas of the US Global Change Research Program, this figure includes only "Improving Knowledge," "Improving Understanding," and "Modelling and Prediction." It also includes about $1 billion of science expenditures under the 2009 Recovery Act that are climate related. The US GAO (2003) puts the 2003 expenditure at about $4 billion.

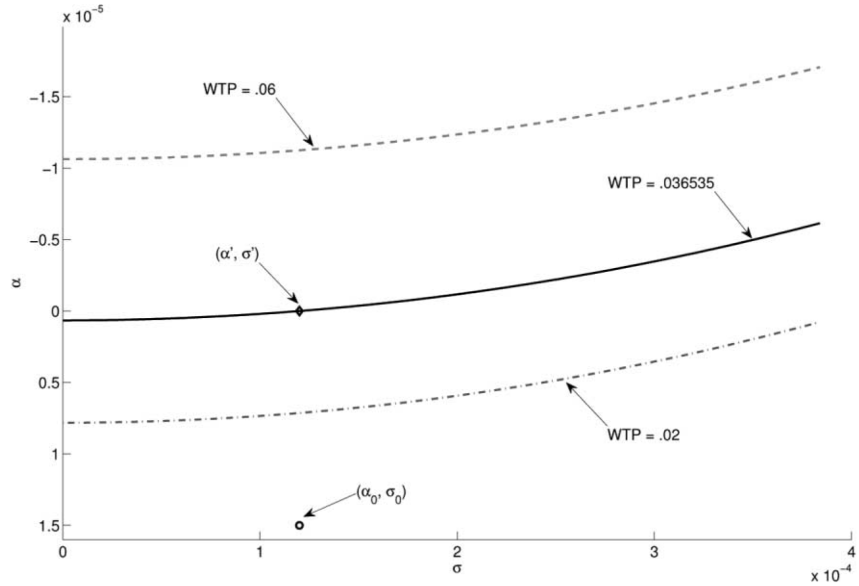

Figure 3. Iso-WTP curves. The starting point is α_0 = .000015 and σ_0 = .00012. The WTP to reach the target (α' = 0, σ' = .00012) is .0365. Moving from (α_0, σ_0) to any other point on the middle curve also has a WTP of .0365. Moving from (α_0, σ_0) to any point on the bottom (top) curve has a WTP of .02 (.06). Also, g_0 = .020, δ = 0, η = 2, and T = 300.

which is shown as a small circle. Also shown is an iso-cost line that is tangent to the iso-WTP curve, that is, at the point where the slope of the iso-WTP curve is $-c_2/c_1$. (Remember that α is decreasing as we move up the vertical axis.) Thus, for this target, $(\alpha', \sigma') = (.0000077, .000046)$, the cost of reaching the iso-WTP curve is minimized. In this example, that cost turns out to be about .020, that is, the same as the WTP. Thus the policy illustrated by figure 5 is just feasible; society would be paying its reservation price for the policy, so the social surplus would be zero. On the other hand, if we were to double c_1 and c_2, the cost would double, making the policy economically infeasible. And if we were to halve c_1 and c_2, the cost would drop to .010, so that the policy would have a social surplus of .02 − .01 = 1% of GDP.

Given equation (19), the cost of any policy is a simple function of the changes in α and σ. But the WTP for changes in α and σ depends not only on the magnitudes of the changes but also on the parameters η, δ, and g_0. For any given shift in α and σ, for example, a reduction in η will increase the WTP.[12] Equivalently, a lower value of η implies that for any fixed WTP, the required changes in α and σ will be smaller. This

12. Reducing η reduces the effective consumption discount rate, which, ignoring uncertainty, is given by $R_t = \delta + \eta g_t$ (the deterministic Ramsey rule). Thus, reducing η increases

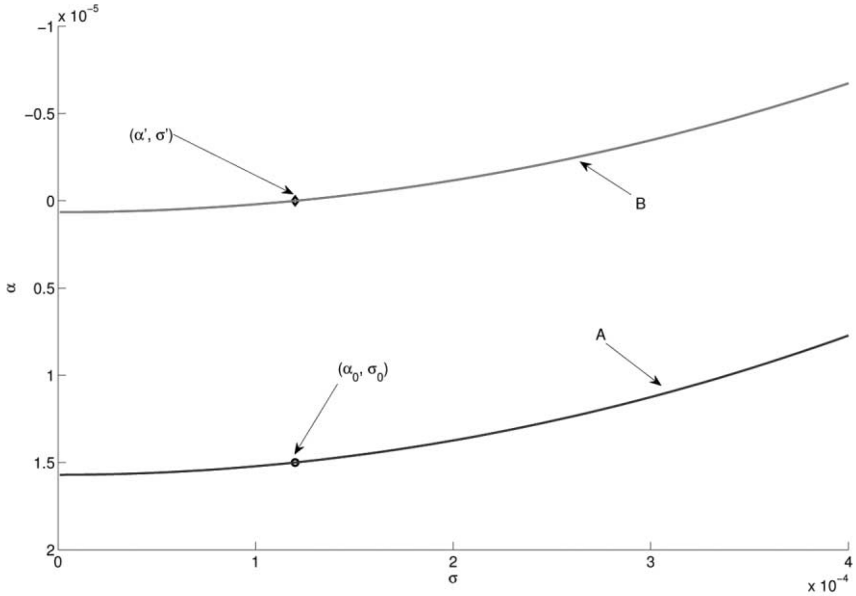

Figure 4. Iso-WTP starting and target curves. Moving from any point on curve A to any point on curve B has a WTP of .0365. (Again, g_0 = .020, $\delta = 0$, $\eta = 2$, T = 300.)

is illustrated by figure 6, which is the same as figure 5 except that $\eta = 1.5$ instead of 2. If $\eta = 1.5$ it takes much smaller reductions in α and σ to have a WTP of .02, and thus the cost of the policy is much smaller (.008 versus .020). In this case the policy has a positive social surplus.

Table 1 shows the target (α', σ'), the cost, and the net social surplus for other parameter values. The first row shows the base case corresponding to figure 5, and the second row corresponds to figure 6 (with η reduced to 1.5 but the WTP fixed at .02). Increasing η to 3 implies that for any fixed target (α', σ'), the WTP will fall, so that obtaining a WTP of .02 requires a much larger shift in (α', σ') and thus a much higher cost (in this case .091, making the social surplus negative). As the fourth row shows, with $\eta = 3$ we can still obtain a positive social surplus, but only if the target (α', σ') has a very small WTP. The fifth row shows that reducing g_0 has the same effect as reducing η (the initial deterministic consumption discount rate is $R_0 = \delta + \eta g_0$), so the target (α', σ') is closer to the initial (α_0, σ_0) and the cost falls. Increasing δ has the same directional effect as increasing η, so with WTP fixed at .02, the cost increases. Finally, in the last row of the table, c_2, the cost of reducing σ is

the present value of future losses of consumption resulting from higher temperatures. For further discussion of this point, see Pindyck (2012b).

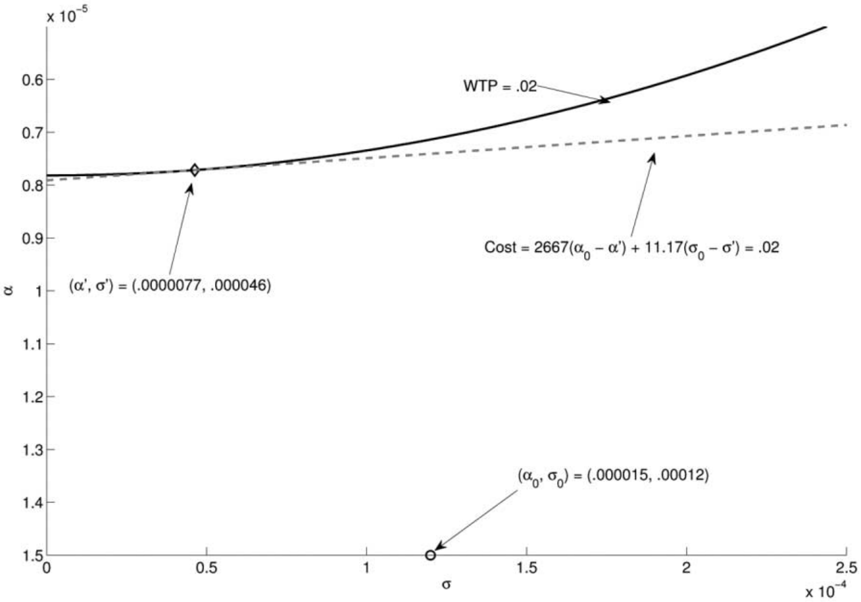

Figure 5. Iso-WTP and iso-cost curves. Moving from α_0 = .000015 and σ_0 = .00012 to any point on the curve has a WTP of .02. The iso-cost line is tangent at (α', σ') = (.0000077, .000046), and total cost is .0203 (i.e., 2.03% of consumption). (g_0 = .020, δ = 0, η = 2, T = 300.)

doubled. This increases the total cost (but not by much), and causes the target (α', σ') to shift to a larger reduction in α and smaller reduction in σ.

2.3. Policy Design

Taking this model at face value, we can infer the optimal mix of risk and return as policy targets. We can also determine if any pair of policy targets (whether or not an optimal mix) yields a positive social surplus, that is, costs less than the WTP. Finally, we can determine alternative target combinations that have the same WTP. In figure 5, for example, the iso-WTP curve corresponds to a WTP of .02, and given the starting point, the optimal target mix cuts both α and σ roughly in half. The cost of reaching this target is .02 as well, so the policy is just feasible. We could have instead chosen an alternative target mix that also has a WTP of .02 (i.e., a different point on the iso-WTP curve), but its cost would exceed .02.

As figure 6 and table 1 show, the social surplus for any policy depends not only on the iso-WTP curve and cost parameters c_1 and c_2, but also on δ, η, and g_0. As discussed in Pindyck (2013a, 2013b), the dependence on δ and η has been a problem for climate policy. What are the "correct" values that should be used for policy evaluation? Economists disagree. The answer depends in part on whether these are behavioral parameters (i.e., reflecting the behavior of consumers, investors, and

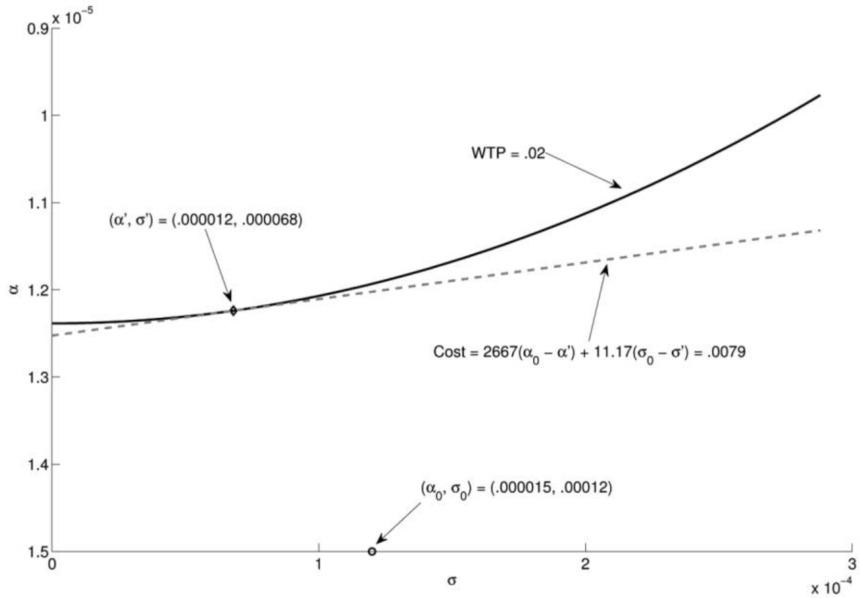

Figure 6. Iso-WTP and iso-cost curves for $\eta = 1.5$. The starting point is again α_0 = .000015 and σ_0 = .00012, but a WTP of .02 now corresponds to larger target values. The iso-cost line is now tangent at (α', σ') = (.000012, .000068), and the total cost is .0079.

Table 1. WTP and Cost

Change from Base Case	WTP	α'	σ'	Cost	Social Surplus
BASE CASE	.02	.0000077	.000046	.020	0
η = 1.5	.02	.000012	.000068	.008	.012
η = 3	.02	−.0000187	.000046	.091	−.071
η = 3	.00032	.0000152	.000041	.00032	0
g_0 = .01	.02	.0000131	.000033	.0060	.014
δ = .01	.02	−.0000031	.000067	.049	−.029
c_2 = 22.34	.02	.0000074	.000090	.021	−.001

Note.—Entries show changes from the base case, for which c_1 = 2,667, c_2 = 11.17, η = 2, g_0 = .02, and δ = 0. Throughout, the starting point is α_0 = .000015 and σ_0 = .00012, but the cost-minimizing target point (α', σ') varies.

firms) or policy parameters (i.e., reflecting the objectives of policy makers). Either way, we have a wide range of reasonable values and thus a wide range of WTP estimates for any target mix. As a comparison of figures 5 and 6 show, a value of η of 1.5 versus 2.0 yields very different WTPs for any set of targets (or, as in the figures, very different targets for a fixed WTP). The approach to policy design outlined in this paper does not get around this problem. Instead it is best understood as a way of clarifying policy target trade-offs (risk versus return) conditional on values for η, and so forth.

3. CONCLUDING REMARKS

I have presented a simple framework by which one can estimate the extent to which the objective of policy should be a reduction in the expected future damages versus a reduction in the uncertainty over those damages. The framework is based on the computation of "iso-WTP curves" (social indifference curves) for combinations of risk and expected returns as policy objectives. Given cost estimates for reducing risk and increasing expected returns, one can compute the optimal risk-return policy mix and evaluate the policy's social surplus.

To illustrate this framework, I used a simple Brownian motion process for the stock variable and a CRRA utility function. This has the advantage of yielding simple expressions for expected future consumption and utility. However, the framework can be applied to alternative stochastic processes for the stock variable and alternative social preferences, although this will introduce additional complexity and numerical methods may be needed to obtain results. Likewise, one might use a more general consumption-based Capital Asset Pricing Model (CAPM) framework in which the discount rate depends on the correlation of benefits (from reducing α and/or σ) with consumption.

For my example of climate change, there is a consensus on the expected rate of temperature increase under BAU—about 3°C. There is considerable disagreement, however, over the nature and extent of uncertainty around that expected value. That disagreement relates to the choice of probability distribution that best applies to future temperatures (e.g., whether it is fat- or thin-tailed), and to the parameterization of any particular distribution.[13] I have used a simple (thin-tailed) normal distribution to describe temperature change, but one can incorporate alternative distributions (using alternative stochastic processes for X_t). Thus, this framework can be used to study the implications of alternative characterizations of uncertainty for determining the targets of policy.

It is an understatement to say that caveats are needed. First, costs of reducing α and σ are probably convex, not linear as in equation (18). But a linear cost function may be a reasonable first approximation, and in any case our ability to estimate these costs is limited. More troublesome is the assumption that whatever the cost, we can change α and σ instantly. One would expect that shifts in α and σ would

13. Weitzman (2009a, 2009b) argues that the distribution for climate change damages is likely to be fat-tailed, which could imply a large "insurance value" for abatement. Pindyck (2011a) shows that a fat-tailed distribution need not imply a high WTP to avert warming and that even if the means and variances are the same, a thin-tailed distribution can lead to a higher WTP.

occur gradually in response to ongoing expenditures on abatement and R&D. Also, I calculated WTP under the restriction that the flows of costs are constant percentages of GDP. In fact, those costs might fall as a result of technological change making it cheaper to limit GHG emissions. And even if the costs do not fall, WTP might be calculated under the alternative assumption that costs increase or decrease over time as the policy becomes more or less stringent. The model can be modified to allow for gradual changes in α and σ and/or costs that are time varying (although both changes would be at the cost of complicating the numerical solution of the model).

Even maintaining the simple assumptions used here (in which the changes in α and σ occur instantly and the flow cost is a constant percentage of GDP), the framework can be extended in a number of directions. For example, the timing of policy implementation can be made endogenous. Suppose a policy is imposed only when temperature reaches a critical threshold (e.g., 3°C). Given a cost for the policy, we could then determine the temperature threshold that maximizes net welfare.

It is important to emphasize what this paper is not. It is not an analysis or evaluation of climate change policy along the lines of the US Interagency Working Group (2010, 2013), which tried to estimate the social cost of carbon (SCC). That study was based on the use of three integrated assessment models that are far more detailed and complex than the extremely simple model I presented here. Likewise, this paper tells us nothing about catastrophic risk, which I have argued elsewhere (Pindyck 2013a, 2013b) is the key to formulating climate change policy. On the other hand, this paper has shown that the SCC has two components that in principle could be decomposed—a component due to expected damages from an additional unit of carbon emissions and a component due to the increased uncertainty over damages that stems from an additional unit of carbon.

APPENDIX
SOME EXPECTATIONS

There are some key expected values that are used throughout this paper. They are derived below, beginning with $\mathcal{E}_s(C_t^{1-\eta})$.

Expected Utility

Denote $F(C, g, s) = \mathcal{E}_s(C_t^{1-\eta})$, for $t > s$. Then F must satisfy the following partial differential equation (the Kolmogorov Forward Equation):

$$dF = \frac{1}{2}\sigma^2 F_{gg} - \alpha F_g + gCF_C + F_s = 0, \tag{A1}$$

with boundary conditions $F(C, g, t) = C_t^{1-\eta}$ and $F(0, g, s) = 0$. Try, and then verify by substitution, a function of the form:

$$F(C, g, s) = C_s^{1-\eta} \exp[(1-\eta)g_s(t-s) + a(s)]. \tag{A2}$$

The partial derivatives are

$$F_g = (1-\eta)(t-s)F,$$

$$F_{gg} = (1-\eta)^2(t-s)^2 F,$$

$$F_C = (1-\eta)F/C,$$

$$F_s = [-(1-\eta)g_s + a'(s)]F.$$

After substituting into equation (A1) and dividing through by F, we have:

$$a'(s) = -\frac{1}{2}\sigma^2(1-\eta)^2(t-s)^2 + \alpha(1-\eta)(t-s) \tag{A3}$$

so that the function $a(s)$ is

$$a(s) = \frac{1}{6}\sigma^2(1-\eta)^2(t-s)^3 - \frac{1}{2}\alpha(1-\eta)(t-s)^2. \tag{A4}$$

Note that if $\eta > 1$ and $\alpha > 0$, $a(s) > 0$.

Thus, starting at time $s = 0$ with $C_0 = 1$,

$$\mathcal{E}_0(C_t^{1-\eta})e^{-\delta t} = e^{-\delta t + (1-\eta)g_0 t - (1/2)\alpha(1-\eta)t^2 + (1/6)\sigma^2(1-\eta)^2 t^3}. \tag{A5}$$

Note that if $\eta > 1$, as t increases, $\mathcal{E}_0(C_t^{1-\eta})$ first decreases and then eventually increases without bound. Thus the integration over time to compute welfare must be done for a finite horizon.

Expected Consumption

We will also need an expression for $\mathcal{E}_0(C_t)$. Following the same steps as above, it is easy to show that if $C_0 = 1$,

$$\mathcal{E}_0(C_t) = e^{g_0 t - (1/2)\alpha t^2 + (1/6)\sigma^2 t^3}. \tag{A6}$$

Note that if $\alpha^2 > 2g_0\sigma^2$, $\mathcal{E}_0(C_t)$ is first increasing in t up to a local maximum, then will decrease to a local minimum, and then increase without bound. But if $\alpha^2 < 2g_0\sigma^2$, then $\mathcal{E}_0(C_t)$ will always be increasing in t.

REFERENCES

Bansal, Ravi, and Marcelo Ochoa. 2011. Welfare costs of long-run temperature shifts. NBER Working Paper 17574 (November). National Bureau of Economic Research, Cambridge, MA.

———. 2012. Temperature, aggregate risk, and expected returns. NBER Working Paper 17575 (January). National Bureau of Economic Research, Cambridge, MA.

Bansal, Ravi, and Amir Yaron. 2004. Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance* 59:1481–1509.

Barrett, Scott. 2008. The incredible economics of geoengineering. *Environmental and Resource Economics* 39:45–54.

———. 2009. The coming global climate-technology revolution. *Journal of Economic Perspectives* 23 (Spring): 53–75.

Dasgupta, Partha. 2008. Discounting climate change. *Journal of Risk and Uncertainty* 37:141–69.

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken. 2009. Temperature and income: Reconciling new cross-sectional and panel estimates. *American Economic Review* 99 (May): 198–204.

———. 2012. Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics* 4:66–95.

Energy Information Administration. 1998. Comparing cost estimates for the Kyoto Protocol. Report no. SR/OIAF/98–03 (October), Energy Information Administration, Washington, DC.

Gollier, Christian. 2013. *Pricing the planet's future: The economics of discounting in an uncertain world*. Princeton, NJ: Princeton University Press.

IPCC (Intergovernmental Panel on Climate Change). 2007a. *Climate change 2007: Impacts, adaptation, and vulnerability*. Cambridge: Cambridge University Press.

———. 2007b. *Climate change 2007: Mitigation of climate change*. Cambridge: Cambridge University Press.

———. 2007c. *Climate change 2007: The physical science basis*. Cambridge: Cambridge University Press.

Kelly, David L., and Charles Kolstad. 1999. Bayesian learning, growth, and pollution. *Journal of Economic Dynamics and Control* 23:491–518.

Lucas, Robert E., Jr. 1987. *Models of business cycles*. Oxford: Blackwell.

Martin, Ian. 2008. Disasters and the welfare cost of uncertainty. *American Economic Review* 98 (May): 74–78.

Nordhaus, William D. 2008. *A question of balance: Weighing the options on global warming policies*. New Haven, CT: Yale University Press.

Palmer, Lisa. 2011. Show me the money. *Nature Climate Change* 1 (October): 376–80.

Pindyck, Robert S. 2007. Uncertainty in environmental economics. *Review of Environmental Economics and Policy* 1 (Winter): 45–65.

———. 2011a. Fat tails, thin tails, and climate change policy. *Review of Environmental Economics and Policy* 5 (Summer): 258–74.

———. 2011b. Modeling the impact of warming in climate change economics. In *The economics of climate change: Adaptations past and present*, ed. G. Libecap and R. Steckel. Chicago: University of Chicago Press.

———. 2012a. Risk and return in environmental economics. NBER Working Paper 18262 (July). National Bureau of Economic Research, Cambridge, MA.

———. 2012b. Uncertain outcomes and climate change policy. *Journal of Environmental Economics and Management* 63 (February): 289–303.

———. 2013a. Climate change policy: What do the models tell us? *Journal of Economic Literature* 51 (September): 860–72.

———. 2013b. The climate policy dilemma. *Review of Environmental Economics and Policy* 7 (Summer): 219–37.

Pindyck, Robert S., and Neng Wang. 2013. The economic and policy consequences of catastrophes. *American Economic Journal: Economic Policy* 5 (November): 306–39.

Samuelson, Paul A. 1956. Social indifference curves. *Quarterly Journal of Economics* 70 (February): 1–22.

Stern, Nicholas. 2007. *The economics of climate change: The Stern Review*. Cambridge: Cambridge University Press.

US Government Accountability Office. 2003. Federal climate change expenditures: Report to Congress. GAO-05-461 (August).

US Interagency Working Group on Social Cost of Carbon. 2010. Technical support document: Social cost of carbon for regulatory impact analysis. Washington, DC: US Government Printing Office.

———. 2013. Technical support document: Technical update of the social cost of carbon for regulatory impact. Washington, DC: US Government Printing Office.

Weitzman, Martin L. 2009a. Additive damages, fat-tailed climate dynamics, and uncertain discounting. *Economics: The Open-Access E-Journal* 3:39.

———. 2009b. On modeling and interpreting the economics of catastrophic climate change. *Review of Economics and Statistics* 91 (February): 1–19.